

Probabilistic Frame Induction*

Jackie Chi Kit Cheung[†]
 Department of Computer Science
 University of Toronto
 Toronto, ON, M5S 3G4, Canada
 jcheung@cs.toronto.edu

Hoifung Poon
 One Microsoft Way
 Microsoft Research
 Redmond, WA 98052, USA
 hoifung@microsoft.com

Lucy Vanderwende
 One Microsoft Way
 Microsoft Research
 Redmond, WA 98052, USA
 lucyv@microsoft.com

Abstract

In natural-language discourse, related events tend to appear near each other to describe a larger scenario. Such structures can be formalized by the notion of a *frame* (a.k.a. template), which comprises a set of related events and prototypical participants and event transitions. Identifying frames is a prerequisite for information extraction and natural language generation, and is usually done manually. Methods for inducing frames have been proposed recently, but they typically use ad hoc procedures and are difficult to diagnose or extend. In this paper, we propose the first probabilistic approach to frame induction, which incorporates frames, events, participants as latent topics and learns those frame and event transitions that best explain the text. The number of frames is inferred by a novel application of a split-merge method from syntactic parsing. In end-to-end evaluations from text to induced frames and extracted facts, our method produced state-of-the-art results while substantially reducing engineering effort.

1 Introduction

Events with causal or temporal relations tend to occur near each other in text. For example, a bombing scenario in an article on terrorism might begin with a DETONATION event, in which terrorists set

off a bomb. Then, a DAMAGE event might ensue to describe the resulting destruction and any casualties, followed by an INVESTIGATION event covering subsequent police investigations. Afterwards, the bombing scenario may transition into a criminal-processing scenario, which begins with police catching the terrorists, and proceeds to a trial, sentencing, etc. A common set of participants serves as the event arguments; e.g., the agent (or subject) of DETONATION is often the same as the theme (or object) of INVESTIGATION and corresponds to the PERPETRATOR.

Such structures can be formally captured by the notion of a *frame* (a.k.a. template), which consists of a set of *events* with prototypical transitions, as well as a set of *slots* representing the common participants. Identifying frames is an explicit or implicit prerequisite for many NLP tasks. Information extraction, for example, stipulates the types of events and slots that are extracted for a frame or template. Online applications such as dialogue systems and personal-assistant applications also model users' goals and subgoals using frame-like representations, and in natural-language generation, frames are often used to represent content to be expressed as well as to support surface realization.

Until recently, frames and related representations have been manually constructed, which has limited their applicability to a relatively small number of domains and a few slots within a domain. Furthermore, additional manual effort is needed after the frames are defined in order to extract frame components from text (e.g., in annotating examples and designing features to train a supervised learning model).

^{*}This is a postprint version of a paper to appear in the *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL/HLT 2013)*.

[†]This research was undertaken during the author's internship at Microsoft Research.

This paradigm makes it hard to generalize across tasks and might suffer from annotator bias.

Recently, there has been increasing interest in automatically inducing frames from text. A notable example is Chambers and Jurafsky (2011), which first clusters related verbs to form frames, and then clusters the verbs’ syntactic arguments to identify slots. While Chambers and Jurafsky (2011) represents a major step forward in frame induction, it is also limited in several aspects. The clustering used ad hoc steps and customized similarity metrics, as well as an additional retrieval step from a large external text corpus for slot generation. This makes it hard to replicate their approach or adapt it to new domains. Lacking a coherent model, it is also difficult to incorporate additional linguistic insights and prior knowledge.

In this paper, we present *ProFinder* (PRObabilistic Frame INDucER), which is the first probabilistic approach for frame induction. *ProFinder* defines a joint distribution over the words in a document and their frame assignments by modeling frame and event transition, correlations among events and slots, and their surface realizations. Given a set of documents, *ProFinder* outputs a set of induced frames with learned parameters, as well as the most probable frame assignments that can be used for event and entity extraction. The numbers of events and slots are dynamically determined by a novel application of the split-merge approach from syntactic parsing (Petrov et al., 2006). In end-to-end evaluations from text to entity extraction using the standard MUC and TAC datasets, *ProFinder* achieved state-of-the-art results while significantly reducing engineering effort and requiring no external data.

2 Related Work

In information extraction and other semantic processing tasks, the dominant paradigm requires two stages of manual effort. First, the target representation is defined manually by domain experts. Then, manual effort is required to construct an extractor or annotate examples to train a machine-learning system. Recently, there has been a burgeoning body of work in alleviating such manual effort. For example, a popular approach to reduce annotation effort is

bootstrapping from seed examples (Patwardhan and Riloff, 2007; Huang and Riloff, 2012). However, this still requires prespecified frames or templates, and selecting seed words is often a challenging task due to semantic drift (Curran et al., 2007). Open IE (Banko and Etzioni, 2008) reduces the manual effort to designing a few domain-independent relation patterns, which can then be applied to extract relational triples from text. While extremely scalable, this approach can only extract atomic factoids within a sentence, and the resulting triples are noisy, non-canonized text fragments.

More relevant to our approach is the recent work in unsupervised semantic induction, such as unsupervised semantic parsing (Poon and Domingos, 2009), unsupervised semantical role labeling (Swier and Stevenson, 2004) and induction (Lang and Lapata, 2011, e.g.), and slot induction from web search logs (Cheung and Li, 2012). As in *ProFinder*, they also model distributional contexts for slot or role induction. However, these approaches focus on semantics in independent sentences, and do not capture discourse-level dependencies.

The modeling component for frame and event transitions in *ProFinder* is similar to a sequential topic model (Gruber et al., 2007), and is inspired by the successful applications of such topic models in summarization (Barzilay and Lee, 2004; Daumé III and Marcu, 2006; Haghighi and Vanderwende, 2009, inter alia). There are, however, two main differences. First, *ProFinder* contains not a single sequential topic model, but two (for frames and events, respectively). In addition, it also models the interdependencies among events, slots, and surface text, which is analogous to the USP model (Poon and Domingos, 2009). *ProFinder* can thus be viewed as a novel combination of state-of-the-art models in unsupervised semantics and discourse modeling.

In terms of aim and capability, *ProFinder* is most similar to Chambers and Jurafsky (2011), which culminated from a series of work for identifying correlated events and arguments in narrative (Chambers and Jurafsky, 2008; Chambers and Jurafsky, 2009). By adopting a probabilistic approach, *ProFinder* has a sound theoretical underpinning, and is easy to modify or extend. For example, in Section 3, we show how *ProFinder* can easily be

augmented with additional linguistically-motivated features. Likewise, `ProFinder` can easily be used as a semi-supervised system if some slot designations and labeled examples are available.

The idea of representing and capturing stereotypical knowledge has a long history in artificial intelligence and psychology, and has assumed various names such as *frames* (Minsky, 1974), *schemata* (Rumelhart, 1975), and *scripts* (Schank and Abelson, 1977). In the linguistics and computational linguistics communities, frame semantics (Fillmore, 1982) uses frames as the central representation of word meaning, culminating in the development of FrameNet (Baker et al., 1998), which contains over 1000 manually annotated frames. A similarly rich lexical resource is the MindNet project (Richardson et al., 1998). Our notion of frame is related to these representations, but there are also subtle differences. For example, Minsky’s frame emphasizes *inheritance*, which we do not model in this paper. (It should be a straightforward extension: using the split-and-merge approach, `ProFinder` already produces a hierarchy of events and slots in learning, although currently, it simply discards the intermediate levels.) As in semantic role labeling, FrameNet focuses on semantic roles and does not model event or frame transitions, so the scope of its frames is often no more than an event in our model. Perhaps the most similar to our frame is Roger Schank’s scripts, which capture prototypical events and participants in a scenario such as restaurant dining. In their approach, however, scripts are manually defined, making it hard to generalize. In this regard, our work may be viewed as an attempt to revive a long tradition in AI and linguistics, by leveraging the recent advances in computational power, NLP, and machine learning.

3 Probabilistic Frame Induction

In this section, we present `ProFinder`, a probabilistic model for frame induction. Let \mathcal{F} be a set of frames, where each frame $F = (E_F, S_F)$ comprises a unique set of events E_F and slots S_F . Given a document D and a word w in D , $Z_w = (f, e)$ represents an assignment of w to frame $f \in \mathcal{F}$ and frame element $e \in E_f \cup S_f$. At the heart of `ProFinder` is a generative model $P_\theta(D, Z)$ that defines a joint

distribution over document D and the frame assignment to its words Z . Given a set of documents \mathcal{D} , frame induction in `ProFinder` amounts to determining the number of frames, events and slots, as well as learning the parameters θ by summing out the latent assignments Z to maximize the likelihood of the document set

$$\prod_{D \in \mathcal{D}} P_\theta(D).$$

The induced frames identify the key event structures in the document set. Additionally, `ProFinder` can also conduct event and entity extraction by computing the most probable frame assignment Z . In the remainder of the section, we first present the base model for `ProFinder`. We then introduce several linguistically motivated refinements, and efficient algorithms for learning and inference in `ProFinder`.

3.1 Base Model

The probabilistic formulation of `ProFinder` makes it extremely flexible for incorporating linguistic intuition and prior knowledge. In this paper, we design our `ProFinder` model to capture three types of dependencies.

Frame transitions between clauses A sentence contains one or more clauses, each of which is a minimal unit expressing a proposition. A clause is unlikely to straddle across different frames, so we stipulate that the words in a clause be assigned to the same frame. On the other hand, frame transitions can happen between clauses, and we adopt the common Markov assumption that the frame of a clause only depends on the clause immediately to its left. Here, sentences are ordered sequentially as they appear in the documents. Clauses are automatically extracted from the dependency parse and further decomposed into an *event head* and its syntactic arguments; see the experiment section for details.

Event transitions within a frame Events tend to transition into related events in the same frame, as determined by their causal or temporal relations. Each clause is assigned an event compatible with its frame assignment (i.e., the event is in the given frame). As for frame transitions, we assume that the event assignment of a clause depends only on the event of the previous clause.

Emission of event heads and slot words Similar to topics in topic models, each event determines a multinomial from which the event head is generated. E.g., a detonation event might use verbs such as *detonate*, *set off* or nouns such as *denotation*, *bombing* as its event head. Additionally, as in USP (Poon and Domingos, 2009), an event also contains a multinomial of slots for each of its argument types¹. E.g., the agent argument of a detonation event is generally the PERPETRATOR slot of the BOMBING frame. Finally, each slot has its own multinomials for generating the argument head and dependency label, regardless of the event.

Formally, let D be a document and C_1, \dots, C_l be its clauses, the ProFinder model is defined by

$$\begin{aligned} P_\theta(D, Z) = & P_{\text{F-INIT}}(F_1) \times \prod_i P_{\text{F-TRAN}}(F_{i+1}|F_i) \\ & \times P_{\text{E-INIT}}(E_1|F_1) \\ & \times \prod_i P_{\text{E-TRAN}}(E_{i+1}|E_i, F_{i+1}, F_i) \\ & \times \prod_i P_{\text{E-HEAD}}(e_i|E_i) \\ & \times \prod_{i,j} P_{\text{SLOT}}(S_{i,j}|E_{i,j}, A_{i,j}) \\ & \times \prod_{i,j} P_{\text{A-HEAD}}(a_{i,j}|S_{i,j}) \\ & \times \prod_{i,j} P_{\text{A-DEP}}(dep_{i,j}|S_{i,j}) \end{aligned}$$

Here, F_i, E_i denote the frame and event assignment to clause C_i , respectively, and e_i denotes the event head. For the j -th argument of clause i , $S_{i,j}$ denotes the slot assignment, $A_{i,j}$ the argument type, $a_{i,j}$ the head word, and $dep_{i,j}$ the dependency from the event head. $P_{\text{E-TRAN}}(E_{i+1}|E_i, F_{i+1}, F_i) = P_{\text{E-INIT}}(E_{i+1}|F_{i+1})$ if $F_{i+1} \neq F_i$.

Essentially, ProFinder combines a frame HMM with an event HMM, where the first models frame transition and emits events, and the second models event transition within a frame and emits argument slots.

¹USP generates the argument types along with events from clustering. For simplicity, in ProFinder we simply classify a syntactic argument into subject, object, and prepositional object, according to its Stanford dependency to the event head.

3.2 Model refinements

The base model captures the main dependencies in event narrative, but it can be easily extended to leverage additional linguistic intuition. ProFinder incorporates three such refinements.

Background frame Event narratives often contain interjections of general content common to all frames. For example, in newswire articles, ATTRIBUTION is commonplace to describe who said or reported a particular quote or fact. To avoid contaminating frames with generic content, we introduce a background frame with its own events, slots, and emission distributions, and a binary switch variable $B_i \in \{BKG, CNT\}$ that determines whether clause i is generated from the actual content frame F_i (CNT) or background (BKG). We also stipulate that if background is chosen, the nominal frame stays the same as the previous clause.

Stickiness in frame and event transitions Prior work has demonstrated that promoting topic coherence in natural-language discourse helps discourse modeling (Barzilay and Lee, 2004). We extend ProFinder to leverage this intuition by incorporating a “stickiness” prior (Haghighi and Vanderwende, 2009) to encourage neighboring clauses to stay in the same frame. Specifically, along with introducing the background frame, the frame transition component now becomes

$$P_{\text{F-TRAN}}(F_{i+1}|F_i, B_{i+1}) = \begin{cases} \mathbf{1}(F_{i+1} = F_i), & \text{if } B_{i+1} = BKG \\ \beta \mathbf{1}(F_{i+1} = F_i) + (1 - \beta) P_{\text{F-TRAN}}(F_{i+1}|F_i), & \text{if } B_{i+1} = CNT \end{cases} \quad (1)$$

where β is the stickiness parameter, and the event transition component correspondingly becomes

$$P_{\text{E-TRAN}}(E_{i+1}|E_i, F_{i+1}, F_i, B_{i+1}) = \begin{cases} \mathbf{1}(E_{i+1} = E_i), & \text{if } B_{i+1} = BKG \\ P_{\text{E-TRAN}}(E_{i+1}|E_i), & \text{if } B_{i+1} = CNT, F_i = F_{i+1} \\ P_{\text{E-INIT}}(E_{i+1}), & \text{if } B_{i+1} = CNT, F_i \neq F_{i+1} \end{cases} \quad (2)$$

Argument dependencies as caseframes As noticed in previous work such as Chambers and Jurafsky (2011), the combination of an event head

syntactic parsing for inferring refined latent syntactic categories (Petrov et al., 2006). Specifically, we initialize our model such that each frame is associated with one event and two slots. Then, after a number of iterations of EM, we split each event and slot in two along with their probability, and duplicate the associated emission distributions. We then add some perturbation to break symmetry. After splitting, we merge back a proportion of the newly split events and slots that result in the least improvement in the likelihood of the training data. For more details on split-merge, see (Petrov et al., 2006)

By adjusting the number of split-merge cycles and the merge parameters, our model learns the number of events and slots in a dynamical fashion that is tailored to the data. Moreover, our model starts with a small number of frame elements, which reduces the number of local optima and make initial learning easier. After each split, the subsequent learning starts with (a perturbed version of) the previously learned parameters, which makes a good initialization that is crucial for EM. Finally, it is also compatible with the hierarchical nature of events and slots. For example, slots can first be coarsely split into persons versus locations, and later refined into subcategories such as perpetrators and victims.

4 MUC-4 Entity Extraction Experiments

We first evaluate our model on a standard entity extraction task, using the evaluation settings from Chambers and Jurafsky (2011) to enable a head-to-head comparison. Specifically, we use the MUC-4 data set (muc, 1992), which contains 1300 training and development documents on terrorism in South America, with 200 additional documents for testing. MUC-4 contains four templates: attack, kidnapping, bombing, and arson.² All templates share the same set of predefined slots, with the evaluation focusing on the following four: perpetrator, physical target, human target, and instrument.

For each slot in a MUC template, the system first identified an induced slot that best maps to it by F_1 on the development set. As in Chambers and Jurafsky (2011), template is ignored in final evaluation. So the system merged the induced slots across all

templates to calculate the final scores. Correctness is determined by matching head words, and slots marked as optional in MUC are ignored when computing recall. All hyper-parameters are tuned on the development set³.

Document classification The MUC-4 dataset contains many documents that contain words related to MUC slots (e.g., *plane* and *aviation*), but are not about terrorism. To reduce precision errors, Chambers and Jurafsky’s (2011) (henceforth, C&J) first filtered irrelevant documents based on the specificity of event heads to learned frames. To estimate the specificity, they used additional data retrieved from a large external corpus. In ProFinder, however, specificity can be easily estimated using the probability distributions learned during training. In particular, we define the probability of an event head in a frame j :

$$P_F(w) = \sum_{E_F \in F} P_{E-HEAD}(w|E)/|F|, \quad (3)$$

and the probability of a frame given an event head:

$$P(F|w) = P_F(w) / \sum_{F' \in \mathcal{F}} P_{F'}(w). \quad (4)$$

We then follow the rest of Chambers and Jurafsky (2011) to score each learned frame with each MUC document, mapping a document to a frame if the average $P_F(w)$ in the document is above a threshold and the document contains at least one *trigger word* w' with $P(F|w') > 0.2$. The threshold and the induced frame were determined on the development set, which were then used to filter irrelevant documents in the test set.

Results Compared to C&J, ProFinder is conceptually much simpler, involving a single probabilistic model, with standard learning and inference algorithms. In particular, it did not require multiple processing steps or customized similarity metrics; rather, it only used the data within MUC-4. In contrast, C&J required additional text to be retrieved from a large external corpus (Gigaword (Graff et al., 2005)) for each event cluster, yet ProFinder nevertheless was able to outperform C&J on entity extraction, as shown in Table 1. Our system achieved

²Two other templates have negligible counts and are ignored as in Chambers and Jurafsky (2011).

³We will make the parameter settings used in all experiments publicly available.

Unsupervised methods	<i>P</i>	<i>R</i>	<i>F₁</i>
ProFinder (This work)	32	37	34
Chambers and Jurafsky (2011)	48	25	33
With extra information			
ProFinder +doc. classification	41	44	43
C&J 2011 +granularity	44	36	40

Table 1: Results on MUC-4 entity extraction. C&J 2011 +granularity refers to their experiment in which they mapped one of their templates to five learned clusters rather than one.

Frame: Terrorism	
Event: Attack <i>report, participate, kidnap, kill, release</i>	Event: Discussion <i>hold, meeting, talk, discuss, investigate</i>
Slot: Perpetrator PERSON/ORG Words: <i>guerrilla, police, source, person, group</i> Caseframes: <i>report>nsubj, kidnap>nsubj, kill>nsubj, participate>nsubj, release>nsubj</i>	Slot: Victim PERSON/ORG Words: <i>people, priest, leader, member, judge</i> Caseframes: <i>kill>doobj, murder>doobj, release>doobj, report>doobj, kidnap>doobj</i>

Figure 2: A partial frame learned by ProFinder from the MUC-4 data set, with the most probable emissions for each event and slot. Labels are assigned by the authors for readability.

good recall but was hurt by the lower precision. We investigated the importance of document classification by only extracting from the gold-standard relevant documents (+doc. classification), which led to a substantial improvement in precision, suggesting possible further improvement by better document classification. Also unlike C&J, our system does not currently make use of coreference information.

Figure 2 shows part of a frame that is learned by ProFinder, including some of the standard MUC slots and events. Our method also finds events not annotated in MUC, such as the discussion event. Other interesting events and slots that we noticed include an arrest event (*call, arrest, express, meet, charge*), a peace agreement slot (*agreement, rights, law, proposal*), and an authorities slot (*police, gov-*

- (a) **Accidents and Natural Disasters:**
 WHAT: what happened
 WHEN: date, time, other temporal markers
 WHERE: physical location
 WHY: reasons for accident/disaster
 WHO_AFFECTED: casualties...
 DAMAGES: ... caused by the disaster
 COUNTERMEASURES: rescue efforts...
- (b) (WHEN *During the night of July 17,*)
 (WHAT *a 23-foot <WHAT tsunami*) *hit the north coast of Papua New Guinea (PNG)>*,
 (WHY *triggered by a 7.0 undersea earthquake in the area*).
- (c) WHEN: *night* WHAT: *tsunami, coast*
 WHY: *earthquake*

Figure 3: An example of (a) a frame from the TAC Guided Summarization task with abbreviated slot descriptions, (b) an annotated TAC contributor, and (c) the entities that are extracted for evaluation.

ernment, force, command). The background frame was able to capture many verbs related to reporting, such as *say, continue, add, believe*, although it missed *report*.

5 Evaluating Frame Induction Using Guided Summarization Templates

One issue with the MUC-4 evaluation is the limited variety of templates and entities that are available. Moreover, this data set was specifically developed for information extraction and questions remain whether our approach can generalize beyond it. We thus conducted a novel evaluation using the TAC guided summarization data set, which contains a wide variety of frames and topics. Our evaluation corresponds to a view of summarization as extracting structured information from the source text, and highlights the connection between summarization and information extraction (White et al., 2001).

Data preparation We use the TAC 2010 guided summarization data set for our experiments (Owczarzak and Dang, 2010). This data set provides templates as defined by the task organizers and contains 46 document clusters in five domains, with each cluster comprising 20 documents on a specific topic. Eight human-written model sum-

maries are provided for each document cluster. As part of the Pyramid evaluation method (Nenkova and Passonneau, 2004), these summaries have been manually segmented and labeled with slots from the corresponding template for each segment (Figure 3)⁴.

We first considered defining the task as extracting entities from the source text, but this annotation is not available in TAC, and pilot studies suggested that it required nontrivial effort to train average users to conduct high-quality annotation reliably. We thus defined our task as extracting entities from the model summaries instead. As mentioned earlier, TAC slot annotation is available for summaries. Furthermore, using the summary text has the advantage that slots that are considered important in the domain naturally appear more frequently, whereas unimportant text is filtered out.

Each span that is labeled by a slot is called a *contributor*. We convert the contributors into a form that is more like the previous MUC evaluation, so that we can fairly compare against previous work like C&J that were designed to extract information into that form. Specifically, we extract the head lemma from all the maximal noun phrases found in the contributor. Like in MUC-4, we count a system-extracted noun phrase as a match if this head word matches and is extracted from the same document (i.e., summary). This process can lead to noise, as the meaning of some contributors depend on a larger phrasal unit than a noun phrase, but this heuristic normalizes the representations of the contributors so that they are amenable to our evaluation. We leave the denoising of this process to future work, and believe it should be feasible by crowdsourcing.

Method and experiments The induced entity clusters are mapped to the TAC slots in the TAC frames according to the best F_1 achieved for each TAC slot. However, one issue is that many TAC slots are more general than the type of slots found in MUC. For example, slots like WHY and COUNTERMEASURES likely correspond to multiple slots at the granularity of MUC. Thus, we map the N -best induced slots to TAC slots rather than the 1-best, for

Systems	1-to-1			5-to-1		
	P	R	F_1	P	R	F_1
ProFinder	24	25	24	21	38	27
C&J	58	6.1	11	50	12	20

Table 2: Results on TAC 2010 entity extraction with N -to-1 mapping for $N = 1$ and $N = 5$. Intermediate values of N produce intermediate results, and are not shown for brevity.

N up to 5. We train ProFinder and a reimplementation of C&J on the 920 full source texts of TAC 2010, and test them on the 368 model summaries. We do not provide C&J’s model with access to external data, in order to create fair comparison conditions to our model. We also eliminate a sentence relevance classification step from C&J, and the document relevance classification step from both models, because all sentences in the summary text are expected to be relevant. We tune C&J’s clustering thresholds and the parameters to our model by two-fold cross validation on the summaries, and assume gold summary classification into the five topic categories defined by TAC.

Results The results on TAC are shown in Table 2. The overall results are poorer than for the MUC-4 task, but this task is harder given the greater diversity in frames and slots to be induced. Like in the previous evaluation, our system is able to outperform C&J in terms of recall and F_1 , but not precision. C&J’s method produces many small clusters, which makes it easy to achieve high precision. The N -to-1 mapping procedure can also be seen to favor their method over ours, many small clusters with high precision can be selected to greatly improve recall, which is indeed the case. However, ProFinder with 1-to-1 mapping outperforms C&J even with 5-to-1 mapping.

6 Conclusion

We have presented the first probabilistic approach to frame induction and shown that it achieves state-of-the-art results on end-to-end entity extraction in standard MUC and TAC data sets. Our model is inspired by recent advances in unsupervised semantic induction and in content modeling in summarization, and is easy to extend. We would like to further

⁴The full set of slots is available at <http://www.nist.gov/tac/2010/Summarization/Guided-Summ.2010.guidelines.html>

investigate frame induction evaluation, for example to evaluate event clustering in addition to the slots and entities.

Acknowledgments

We would like to thank Nate Chambers for answering questions about his system. We would also like to thank Chris Quirk for help with preprocessing the MUC corpus, and the other members of the NLP group at Microsoft Research for useful discussions.

References

- [Baker et al.1998] Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *Proceedings of the 17th International Conference on Computational linguistics*.
- [Banko and Etzioni2008] Michele Banko and Oren Etzioni. 2008. The tradeoffs between open and traditional relation extraction. *Proceedings of ACL-08: HLT*, pages 28–36.
- [Barzilay and Lee2004] Regina Barzilay and Lillian Lee. 2004. Catching the drift: Probabilistic content models, with applications to generation and summarization. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*.
- [Bean and Riloff2004] David Bean and Ellen Riloff. 2004. Unsupervised learning of contextual role knowledge for coreference resolution. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*.
- [Chambers and Jurafsky2008] Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio, June. Association for Computational Linguistics.
- [Chambers and Jurafsky2009] Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*. Association for Computational Linguistics.
- [Chambers and Jurafsky2011] Nathanael Chambers and Dan Jurafsky. 2011. Template-based information extraction without the templates. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 976–986, Portland, Oregon, USA, June. Association for Computational Linguistics.
- [Cheung and Li2012] Jackie C. K. Cheung and Xiao Li. 2012. Sequence clustering and labeling for unsupervised query intent discovery. In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining*, pages 383–392.
- [Curran et al.2007] James R. Curran, Tara Murphy, and Bernhard Scholz. 2007. Minimising semantic drift with mutual exclusion bootstrapping. In *Proceedings of the 10th Conference of the Pacific Association for Computational Linguistics*.
- [Daumé III and Marcu2006] Hal Daumé III and Daniel Marcu. 2006. Bayesian Query-Focused summarization. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 305–312, Sydney, Australia, July. Association for Computational Linguistics.
- [Fillmore1982] Charles J. Fillmore. 1982. Frame semantics. *Linguistics in the Morning Calm*, pages 111–137.
- [Graff et al.2005] David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2005. English gigaword second edition. *Linguistic Data Consortium, Philadelphia*.
- [Gruber et al.2007] Amit Gruber, Michael Rosen-Zvi, and Yair Weiss. 2007. Hidden topic markov models. *Artificial Intelligence and Statistics (AISTATS)*.
- [Haghighi and Vanderwende2009] Aria Haghighi and Lucy Vanderwende. 2009. Exploring content models for multi-document summarization. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 362–370, Boulder, Colorado, June. Association for Computational Linguistics.
- [Huang and Riloff2012] Ruihong Huang and Ellen Riloff. 2012. Bootstrapped training of event extraction classifiers. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 286–295, Avignon, France, April. Association for Computational Linguistics.
- [Lang and Lapata2011] Joel Lang and Mirella Lapata. 2011. Unsupervised semantic role induction via split-merge clustering. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1117–1126, Portland, Oregon, USA, June. Association for Computational Linguistics.
- [Liang and Klein2009] Percy Liang and Dan Klein. 2009. Online EM for unsupervised models. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 611–

- 619, Boulder, Colorado, June. Association for Computational Linguistics.
- [Minsky1974] Marvin Minsky. 1974. A framework for representing knowledge. Technical report, Cambridge, MA, USA.
- [muc1992] 1992. *Proceedings of the Fourth Message Understanding Conference (MUC-4)*. Morgan Kaufmann.
- [Nenkova and Passonneau2004] Ani Nenkova and Rebecca Passonneau. 2004. Evaluating content selection in summarization: The pyramid method. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, volume 2004, pages 145–152.
- [Owczarzak and Dang2010] Karolina Owczarzak and Hoa T. Dang. 2010. TAC 2010 guided summarization task guidelines.
- [Patwardhan and Riloff2007] Siddharth Patwardhan and Ellen Riloff. 2007. Effective information extraction with semantic affinity patterns and relevant regions. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 717–727, Prague, Czech Republic, June. Association for Computational Linguistics.
- [Petrov et al.2006] Slav Petrov, Leon Barrett, Romain Thibaux, and Dan Klein. 2006. Learning accurate, compact, and interpretable tree annotation. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*.
- [Poon and Domingos2009] Hoifung Poon and Pedro Domingos. 2009. Unsupervised semantic parsing. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 1–10.
- [Richardson et al.1998] Stephen D. Richardson, William B. Dolan, and Lucy Vanderwende. 1998. MindNet: Acquiring and structuring semantic information from text. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 2*, pages 1098–1102, Montreal, Quebec, Canada, August. Association for Computational Linguistics.
- [Rumelhart1975] David Rumelhart, 1975. *Notes on a schema for stories*, pages 211–236. Academic Press, Inc.
- [Schank and Abelson1977] Roger C. Schank and Robert P. Abelson. 1977. *Scripts, Plans, Goals, and Understanding: An Inquiry Into Human Knowledge Structures*. Lawrence Erlbaum, July.
- [Swier and Stevenson2004] Robert S. Swier and Suzanne Stevenson. 2004. Unsupervised semantic role labelling. In Dekang Lin and Dekai Wu, editors, *Proceedings of EMNLP 2004*, pages 95–102, Barcelona, Spain, July. Association for Computational Linguistics.
- [White et al.2001] Michael White, Tanya Korelsky, Claire Cardie, Vincent Ng, David Pierce, and Kiri Wagstaff. 2001. Multidocument summarization via information extraction. In *Proceedings of the First International Conference on Human Language Technology Research*. Association for Computational Linguistics.